

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank) 2. REPORT DATE July 22, 1996 3. REPORT TYPE AND DATES COVERED Final technical report 3-1-95 / 2-29-96

4. TITLE AND SUBTITLE
Optimization and Artificial Intelligence

5. FUNDING NUMBERS
F49620-95-1-0233

6. AUTHOR(S)
E. Boros, P.L. Hammer and F.S. Roberts

AFOSR TR 96
G400

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
RUTCOR - Rutgers University Center for Operations Research
Rutgers University
P.O. Box 5062
New Brunswick, NJ 08903

8. PERFORMING ORGANIZATION
REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
110 Duncan Avenue, Suite 100
Bolling Air Force Base
Washington, DC 20332-0001

10. SPONSORING/MONITORING
AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

no limitations

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

We propose to do research in a number of areas at the interface between operations research and artificial intelligence: pattern finding and sequential diagnosis, managing data and knowledge bases efficiently, and scheduling problems.

19960729 019

DTIC QUALITY INSPECTED 3

14. SUBJECT TERMS
scheduling, sequential diagnosis, logical analysis, Boolean functions

15. NUMBER OF PAGES
21

16. PRICE CODE

17. SECURITY CLASSIFICATION
OF REPORT
unclassified

18. SECURITY CLASSIFICATION
OF THIS PAGE
unclassified

19. SECURITY CLASSIFICATION
OF ABSTRACT
unclassified

20. LIMITATION OF ABSTRACT
UL



RUTCOR • Rutgers Center for Operations Research
P.O. Box 5062 • New Brunswick • New Jersey 08903-5062
908/445-4856 • FAX: 908/445-5472

RUTCOR

FINAL TECHNICAL REPORT

TO: AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

OPTIMIZATION AND ARTIFICIAL INTELLIGENCE

GRANT NUMBER F49620-95-1-0233

ACCOMPLISHMENTS: March 1, 1995 – February 29, 1996

Endre Boros, co-Principal Investigator

Peter L. Hammer, co-Principal Investigator

Fred S. Roberts, co-Principal Investigator

July 20, 1996

RUTCOR

FINAL TECHNICAL REPORT

TO: AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

OPTIMIZATION AND ARTIFICIAL INTELLIGENCE

GRANT NUMBER F49620-95-1-0233

This summary of research accomplishments is organized into essentially the same sections as is our original proposal. Papers referred to by number are listed below in the list of publications prepared under the grant. Papers referred to with authors' names and year are listed at the end of the section.

1. Overview of the Project and Applied Motivation

In this project, we have done research in a number of areas at the interface between operations research and artificial intelligence. The research problems considered were motivated by a number of practical problems which are of interest and importance to the Air Force. The Air Force problems which motivated the project involve decision support systems, manpower planning and training, scheduling and deliberate planning, pattern detection, and the need to understand trade-offs in decisions with multiple objectives.

Advanced decision support and expert systems are in increasing demand in our highly technological society. Increasingly, the knowledge and data bases used to support decisions at the Air Force and elsewhere are extremely large. Already, it is estimated by the Command Analysis Group at the Air Mobility Command Headquarters [1994] that 50 to 80% of the entire effort for one of their studies is spent in data gathering, validation, and verification. Because of the increasingly large size of knowledge and data base systems used in decision support, the "exhaustive" methods used to gather information and still widely in use are no longer viable. There is clearly a need for the development of powerful algorithms and heuristics for handling large knowledge and data bases, for decomposing them or organizing them in some optimal way, and for using them to detect patterns, make diagnoses and inferences, plan schedules, or compare complex alternatives. The desire to make optimal use of knowledge and data bases in decision support is a major part of the increasingly important and fruitful interface between artificial intelligence (AI) and operations research (OR). Researchers in AI have been trying to develop aids for decisionmaking that mimic intelligent human behavior by being able to organize data, process it rapidly and efficiently, learn from the past, and pinpoint new ideas. Operations researchers have, for years, been successfully modelling practical problems as optimization problems and developing methods for solving these problems. In this project, we have explored the use of optimization methods to solve problems of AI. We have also explored a number of optimization problems from the point of view of the interface between OR and AI and the need to provide efficient decision support for solving optimization problems of the Air Force.

Increasingly, decision makers are faced with a large amount of data, often incomplete or subject to noise, and are asked to detect patterns and make diagnoses and inferences based on these data. The pattern finding - diagnosis -

inference problem arises in a wide variety of contexts of importance to the Air Force. It has long been a central problem of AI, which is concerned with detecting patterns in an environment in order to maneuver through it or find diagnoses, with sequential improvement in performance. Recently, the problem has been approached from the point of view of OR, and optimization methods have been used to find the best inference or diagnosis, the best question to ask in order to improve the possibility of making a good inference, etc. We have explored this interface between AI and OR in some detail.

Scheduling problems are a fundamental part of many activities of the Air Force. There is a long tradition of analysis of such problems in OR, where optimization methods have been used to find efficient schedules for many years. In recent years, researchers in AI have concentrated on scheduling-type problems in design of machines that mimic intelligent search behavior. The optimization methods of OR are increasingly of interest in AI. We have explored issues of scheduling and in particular issues of scheduling at the interface between AI and OR.

As pointed out by the Command Analysis Group at Air Mobility Command Headquarters [1994], "almost every important decision at AMC requires making trade-offs among many objectives or measures of merit." Decisionmaking when there are multiple objectives has been a subject of study by operations researchers for many years. At the heart of this topic is the solution of the optimization problems that can be made precise after explicit representation and analysis of preferences and choices, a representation and analysis that until recently has usually been missing in decision support systems and in models of intelligent behavior. However, things have been changing, and there has recently been increasing interest in the explicit representation of preferences and choices among multiattributed alternatives in the literature of AI. We have briefly explored problems of trade-offs between multiple objectives from both the OR and AI point of view.

While the problems we have studied all lie at the interface between OR and AI, we also feel that many of these problems in applied operations research are of considerable interest in their own right, in particular from the point of view of their Air Force applications, and we have explored them for their own sake in this project. In particular, we have emphasized research on some fundamental questions of scheduling and cluster analysis that are of interest for their many Air Force applications.

The rest of this report is divided into three research sections, in each of which we describe a particular set of mathematical problems and relate them to the motivating applied problems. These sections deal with pattern finding, managing data and knowledge bases efficiently, and connections to scheduling problems and the handling of trade-offs among multiple objectives.

2. Pattern Finding

One of the most difficult of human behaviors to understand, and one of the most difficult to mimic with machines, is inductive reasoning. Humans are remarkably good at inferring general rules from specific instances. We have been concerned with this problem of inference, and specifically the problem of how to infer patterns or explanatory theories based on partial information.

Our interest in this problem was motivated by a number of specific problems. For example, suppose we wish to build a rule base for an expert system for predicting failure of an (electronic or other) system (or network). We have records of past failures of the system when certain of its components fail and when certain others do not. However, we do not have records that cover every possible situation. Thus, the data are *incomplete* or *partial*. Moreover, the failures of the system may not always take place given the failure of certain components. Hence, the data are *noisy*. We would like to derive a rule that predicts, given partial and/or noisy data, whether or not the system will fail in any conceivable instance.

A similar problem arises if we are trying to diagnose a patient who has certain symptoms and not certain others. We may want to determine which combinations of foods cause a patient's suspected food allergy; we ask the patient to record, each day, whether he eats certain foods and whether an allergic reaction develops; and we wish to systematically relate the occurrence of a reaction to which foods were eaten.

To give another example, consider the problem of processing enlistment applications for Air Force officers. We have past records of enlistment officers' decisions. We would like to use these as a basis for formulating a rule that indicates when an enlistment application should be accepted. Each application is characterized by the presence or absence of a fixed set of attributes and by the enlistment officer's decision. Again, the data could be partial and it could be noisy. We wish to build an expert system that will make enlistment decisions automatically, even in situations not covered by previous data. We hope that such an expert system will make decisions more systematically, more rapidly, and more consistently than varying enlistment officers. (A similar problem is faced, for example, by a bank trying to build an expert system for action on loan applications.)

A variant on this problem arises if we try to understand why pilots leave the Air Force to work for commercial airlines. Suppose we have a database of resignations and for each pilot who resigns, we record the presence or absence of certain attributes, such as an advanced degree, ten years or more of service, at least two children, etc. Can we use the pattern of presence or absence of attributes to predict whether or not a given pilot will resign, even if that pilot's set of attributes is not one we have seen before?

To give still another example, consider what happens when we want to teach a robot to maneuver in an area filled with obstacles. An obstacle might appear as a certain pattern of light or dark pixels. In some situations, the pattern of pixels corresponds to an object, in others it does not. The robot must be able to learn from the previously observed data and determine whether or not a new pattern corresponds to an obstacle.

Similar problems arise in "troubleshooting" in many complex systems including networks and electronic and mechanical systems, in searching and sorting in hazardous or nuclear or chemically toxic environments, in detecting enemy positions, in remote operations in space or underseas, and so on.

When only partial observations are available, no method can definitively predict the result that will occur under all possible observations, even if there is no noise present. However, the cause-effect relationship can be narrowed down

sufficiently to provide substantial guidance to a decisionmaker, and it has been our goal to develop methods to do so.

We have taken a Boolean approach, modelling the problem as one of trying to discover a partially defined Boolean function. Models related to the ones we shall build have been considered in the artificial intelligence literature, mostly by researchers interested in machine learning and in inductive inference. Our approach opens the possibility of taking direct advantage, in an AI framework, of an enormous body of known results concerning Boolean functions.

We have modeled the inductive inference problem as a "cause-effect" problem in which there are n identified potentially relevant *Boolean variables* x_1, x_2, \dots, x_n , i.e., variables which can be interpreted to either be present (1) or absent (0) in any given instance. In the following, we shall sometimes call these variables and their negations $\neg x_1, \neg x_2, \dots, \neg x_n$, *positive* and *negative literals*, respectively, or simply *literals*. In the simplest model of our problem, we think of having an unknown *Boolean function* $f(x_1, x_2, \dots, x_n)$, i.e., a function taking on the value 0 or 1 for every combination of Boolean arguments. Suppose that we have observed whether an event occurs (1) or doesn't (0) in certain situations in which we know some of the potentially relevant variables are present and some others are not, and perhaps for some we are not sure. We then say that the Boolean function is *partially defined*. For instance, suppose we know that $f(1,0,x) = 1$ regardless of the value of x and that $f(1,1,1) = 0$, but otherwise we do not know f . Our goal is to predict the value of the Boolean function for any combination of the values of the arguments. For instance, in the theory of network reliability, the variables x_i correspond to different links, with $x_i = 0$ corresponding to the event of the i th edge failing, and the Boolean function taking on the value 0 if the entire network fails.

In the paper [10], we address the fundamental problem of finding a Boolean function (extension) f given a set of data, represented as a set of binary "true n -vectors" (or "positive examples") and a set of "false n -vectors" (or "negative examples"). We seek an extension f with some specified properties so that f is true (respectively false) in every given true (respectively false) vector. We study this problem in the presence of some a priori knowledge about the extension f . Such knowledge may be obtained from experience or from the analysis of mechanisms that may or may not cause the phenomena under consideration. The real-world data may contain errors, e.g., measurement errors might come in when obtaining data, or there may be some other influential factors not represented as variables in the vectors. To cope with such situations, we may have to give up the goal of establishing an extension that is perfectly consistent with the given data. If there is no such extension, the best we can expect is to establish an extension f which has the minimum number of misclassifications. Both problems, i.e., the problem of finding an extension within a specific class of Boolean functions and the problem of finding a minimum error extension in that class, are extensively studied in paper [10], which was begun under an earlier AFOSR project and significantly revised under this one. For certain classes, we provide polynomial algorithms, and for others, we prove their NP-hardness.

As a form of knowledge acquisition from data, we consider in papers [11] and [9] the problem of deciding whether there exists an extension of a partially defined Boolean function with missing data, but with sets of positive and negative examples given. We define three types of extensions, called consistent, robust, and

most robust, depending upon how to deal with missing bits. We study these types of extensions for various classes of Boolean functions, including general, positive, regular, k -DNF, h -term, DNF, Horn, self-dual, threshold, read-once, and decomposable. For certain classes, we provide polynomial time algorithms, while for others we prove NP-hardness.

In paper [6], begun under an earlier AFOSR project and considerably revised under the current one, we consider the problem of identifying an unknown Boolean function f by asking an oracle the functional values $f(a)$ for a selected set of test vectors a in $\{0,1\}^n$. Furthermore, we assume that f is a positive (or monotone) function of n variables. It is not known yet whether or not the whole task of generating test vectors and checking if the identification is completed can be carried out in polynomial time in n and m , where $m = |\min T(f)| + |\max F(f)|$ and $\min T(f)$ (respectively $\max F(f)$) denotes the set of minimal true (respectively, maximal false) vectors of f . To partially answer this question, we propose in this paper two polynomial time algorithms that, given an unknown positive function f of n variables, decide whether or not f is 2-monotonic, and if so, output both sets $\min T(f)$ and $\max F(f)$. The first algorithm uses $O(nm^2 + n^2m)$ time and $O(nm)$ queries, while the second uses $O(n^3m)$ time and $O(n^3m)$ queries.

3. Managing Knowledge and Data Bases Efficiently

As we noted in Section 1, the knowledge and data bases used to support decisions at the Air Force and elsewhere are extremely large and the widely-used methods of gathering information from them are becoming less and less viable. Motivated by the need to manage knowledge and data bases efficiently, we have devoted considerable emphasis in this project to three problems, knowledge base compression through "logic minimization," database decomposition, and knowledge discovery through the generation of previously unknown and potentially useful conclusions from data.

3.1. Logic Minimization

The growing complexity of knowledge incorporated in modern expert systems has led to a rapid increase in the size of their knowledge bases. One consequence of this development is the increase of the response time, due to the dependence of the computational complexity of answering queries on the size of the knowledge base. Knowledge base compression reduces memory requirements and accelerates answering queries, thus leading to a drastic speedup in overall computational performance of expert systems. The importance of knowledge base compression for Air Force applications was underlined in the document by the Command Analysis Group of the Air Mobility Command [1994], where, as we mentioned in Section 1, it was noted how large a percentage of the entire effort for its studies is usually spent in data gathering, validation, and verification. In particular, a scenario being simulated is typically very rich in information and detail, and redundancy needs to be eliminated or minimized, for example to facilitate consistency checking, to make decisions more efficiently, and to generate new conclusions. Redundancy minimization is exactly the problem we have addressed.

The problem of knowledge compression in expert systems has been formalized in this project as a logic minimization problem. To explain this idea, we note that a particular knowledge base is just one possible representation of knowledge, interpreted as the set of models of the knowledge base. It was observed by

Hammer and Kogan [1993] that other, logically equivalent, representations of the same knowledge exist, and they can have significantly smaller size. The problem of *logic minimization* is concerned with finding optimally smaller representations of information in a knowledge base while preserving the set of satisfying models.

Recall that a disjunctive normal form or DNF for a Boolean function is a disjunction of terms, which are conjunctions of literals in which each literal appears at most once. Similarly, a *conjunctive normal form* or *CNF* for a Boolean function is a conjunction of *clauses*, which are disjunctions of literals in which each literal appears at most once. It is well known that every Boolean function can be represented not only by a DNF, but also by a CNF. A CNF of a Boolean function is not unique and any two CNF's of the same Boolean function are called *equivalent*.

We have formulated the problem of logic minimization as the problem of constructing a DNF or a CNF that represents a given Boolean function and is optimal with respect to a certain complexity measure. The two most commonly used complexity measures are the number of terms (clauses) and the *length* (number of literals) in the DNF (CNF). The general problem of finding an optimal representing DNF or CNF is known to be very difficult computationally (NP-hard), which is the reason that most real systems employ approximative algorithms for its solution. A particular type of logic minimization problem may vary from application to application. In the field of high-level synthesis Boolean functions are usually represented by DNFs which are obtained as a result of compiling the initial specification of the problem written in one of a number of hardware description languages. In the field of artificial intelligence, the production rule knowledge bases are in fact sets of clauses (CNFs).

The key features of logical analysis of data are the discovery of minimal sets of features necessary for explaining all observations, and the detection of hidden patterns in the data capable of distinguishing observations describing "positive" outcome events from "negative" outcome events. Combinations of such patterns are used for developing general classification procedures. The paper [7] gives a broad introduction to the topic of logical analysis of data, with an emphasis on numerical data.

In the paper [8], we describe a new, logic-based methodology for analyzing observations based on the general principles about logical analysis of data described in the previous paragraph. An implementation of this methodology is described in the paper, along with the results of numerical experiments demonstrating the classification performance of logical analysis of data in comparison with the reported results of other procedures. In the final section, we describe three pilot studies on applications of logical analysis of data to oil exploration, psychometric testing, and the analysis of developments in the Chinese transitional economy. These pilot studies demonstrate not only the classification power of logical analysis of data, but also its flexibility and capability to provide solutions to various cross-dependent problems.

Paper [15] goes into considerable detail about the logical analysis of economic data about China. It demonstrates how logical analysis allows the development of a decision support system.

In paper [1], we analyze the generalization accuracy of standard techniques for the logical analysis of data, using a probabilistic framework.

3.2. Decompositions of Boolean Functions

Decompositions of a fully or partially defined Boolean function help to save storage and to speed up future queries, and so make management of knowledge and data bases more efficient. The problem of detecting decomposability and finding decompositions has arisen recently at the convergence of ideas from logic minimization, computational complexity, machine learning theory, and image processing in a promising new approach to robust pattern finding and efficient knowledge and data base management. The decomposition approach is described by Dechter and Pearl [1992] and Maier [1983] and also in the papers by Ross, Noviskey, Axtell, and Breen [1993] and Ross, Axtell, Noviskey, and Gadd [1993], which among other things describe work at Wright Laboratory at Wright Paterson Air Force Base. We have investigated several problems related to decomposition.

In paper [13], we study the important class of Horn functions and provide a simple characterization. We then study in detail the special class of submodular functions. We give a one-to-one correspondence between submodular functions and partial preorders (reflexive and transitive binary relations), and in particular between the nondegenerate acyclic submodular functions and the partially ordered sets. This led us to graph-theoretic characterizations of all minimum DNF representations of a submodular function and to show that the problem of recognizing submodular functions in DNF representation is in Co-NP.

Paper [12] is concerned with the variable deletion control set problem, the problem of finding a minimum cardinality set of variables whose deletion from the formula results in a DNF satisfying some prescribed property. Similar problems can be defined with respect to the fixation of variables or the deletion of terms in a DNF. In this paper, begun under an earlier AFOSR-project and revised under this one, we investigate the complexity of such problems for a broad class of DNF properties.

3.3. Knowledge Discovery in Databases and Some Underlying Clustering Problems

One of the main problems we are facing in the information age is that human abilities cannot handle the growth in size and number of existing databases. In order to cope with this information flood, we need to develop specialized tools that will automatically analyze an existing database and generate interesting conclusions based on the information stored in these databases. The research area that has arisen from this phenomenon, that of knowledge discovery in databases, is one of the fastest growing research fields in AI, and it can be expected to have significant impact as we approach the next century. *Knowledge discovery* is defined as the nontrivial extraction of implicit, previously unknown, and potentially useful information from given data (Piatetsky-Shapiro and Frawley [1991]).

Our approach to knowledge discovery was based on a framework developed by Martin Golumbic and Ronen Feldman of Bar-Ilan University in Israel. This process involves three stages, parsing, clustering, and drawing inferences. Each stage requires major theoretical developments at the interface between OR and AI, and we have worked on some of these issues, with an emphasis on clustering.

There are many issues of clustering relevant to knowledge discovery, and we have investigated a variety of them. Clustering methods are not only relevant to

knowledge discovery, but also to the analysis of many practical problems of the Air Force which involve large amounts of data. These problems arise in such diverse contexts as early warning systems, detection of enemy positions, remote operations in space, cargo movement, "troubleshooting" in complex electronic systems, and forecasting. The data that arises is often noisy and unreliable, sometimes arising in hazardous environments, or under jamming, or just subject to great uncertainties. We can use clustering methods to detect patterns or to identify underlying causes. Clustering methods have been used at AMC in solving location problems, for instance in locating (through the OADS model) U.S. hubs at Travis Air Force Base in California and Tinker Air Force Base in Oklahoma, in identifying good points of embarkation in deliberate planning models, in identifying staging areas for medical evacuations, and in identifying hubs for the defense courier system. While developing clustering methods for knowledge discovery, we have kept other Air Force applications of these methods in mind as motivation.

Given a set of points in Euclidean space, and a partitioning of this "training set" into two or more subsets ("classes"), we consider in paper [14] the problem of identifying a "reasonable" assignment of another point in the Euclidean space ("query point") to one of these classes. The various classifications proposed in this paper are determined by the distances between the query point and the points in the training set. We report results of extensive computational experiments comparing the new methods with two well-known distance-based classification methods (k-nearest neighbors and Parzen windows) on data sets commonly used by the machine learning community. The results show that the performance of both new and old distance-based methods is on a par with and often better than that of the other best classification methods known. Moreover, the new classification procedures proposed in this paper are easy to implement, extremely fast, and very robust in the sense that their performance is insignificantly affected by the choice of parameter values.

Paper [16] introduces measures of relevance for sets of variables in a classification knowledge base. Sets of variables which determine the outcome of a classification regardless of the values of the other variables have relevance 1. More generally, the relevance of a set of variables measures the expected degree of certainty of a classification when the values of the variables in the set are known. Properties of a class of relevance-type measure are studied. It is shown that the relevance of a set of variables is not less than that of any of its subsets. Cases of extreme relevance value are characterized. The relationship of relevance and the classic concept of "strength" of a Boolean variable is investigated, and it is proved that sets of stronger variables have higher relevance.

In approaches to clustering where we derive the solution from judgements of closeness or similarity, the interval graph model seems particularly relevant. Here, we start with judgements of closeness, assign to each element being judged a real interval, and take two intervals to overlap if and only if the corresponding elements are close; the real intervals are then used to define the clusters. All of this can be accomplished if and only if the graph whose vertices are the elements and whose edges correspond to closeness defines an *interval graph*. Interval graphs are part of the more general class of graphs called perfect graphs that has a wide variety of important practical applications, including applications to clustering and scheduling problems of various kinds. Our investigation of interval graphs and perfect graphs has led to two papers.

A *matrix* of 0's and 1's is called *perfect* if the associated set packing polytope $P(A) = \{x: Ax \leq 1, 0 \leq x \leq 1\}$ is integral. Perfect matrices have many interesting properties and the perfectness of a 0,1 matrix is closely related to the perfectness of an associated graph. A *matrix* of 0's, 1's, and -1's is called *perfect* if the corresponding generalized set packing polytope $P(A) = \{x: Ax \leq 1 - n(A), 0 \leq x \leq 1\}$ is integral, where $n(A)$ is the vector whose r th component is the number of negative entries in row r of A . In paper [3], begun under an earlier AFOSR project and revised in the present one, we provide a characterization of such perfect matrices in terms of an associated graph which one can build in $O(n^2m)$ time, where $m \times n$ is the size of the matrix. We also obtain an algorithm of the same time complexity, for testing the irreducibility of the corresponding generalized set packing polytope.

We have applied the theory of perfect graphs and hypergraphs to some problems of game theory that are also relevant to design of efficient and reliable networks and to the analysis of multi-attributed utility data. A game can be defined by the set I of players and the set A of outcomes, and a *coalition* is then a subset of I . The *core* of a game is defined as the set of outcomes acceptable for all coalitions and it is probably the simplest and most natural concept of cooperative game theory. In paper [5], we note that some players may not like or know each other, so they cannot form a coalition. Let K be a fixed family of coalitions. The *K-core* is defined as the set of outcomes acceptable for all coalitions from K . The family K is called *stable* if the *K-core* is not empty for any normal form game. We prove that a family K of coalitions is stable if and only if K is a normal hypergraph.

4. Scheduling Problems and Methods for Handling Trade-offs among Multiple Objectives.

Scheduling is one of the basic tasks in almost all planning systems. Many Air Force activities involve *scheduling problems*. For instance, at AMC, scheduling problems arise in allocating loads to airplanes, assigning loads to points of embarkation and to routes, assigning crews to airplanes, and so on. Scheduling theory has long been a major area of interest in operations research, and there have been hundreds of papers written in the field. More recently, there has been a surge of interest in the interface between the scheduling problems of OR and the scheduling techniques of AI. For example, De [1988] presents a knowledge-based approach to scheduling and Feldman and Golumbic [1990] use constraint-satisfiability algorithms to solve scheduling problems. The particular Air Force scheduling problems mentioned above have numerous complications which scheduling theory has not addressed. We have investigated a variety of approaches to scheduling which take into account extra complications motivated by Air Force problems, in particular taking account of user preferences for schedules, finding schedules that meet performance standards such as those embodied in the UMMIPS priorities at AMC, and taking account of conflicting requests for schedules.

4.1. Scheduling Under Performance Constraints

Often the "measures of merit" (such as those mentioned by the Command Analysis Group at Air Mobility Command Headquarters) that are used in multiobjective decision problems, and in particular in scheduling problems, are based on subjective judgements or on scaling procedures that are subject to modification. In solving decisionmaking problems, human decisionmakers often use scales of measurement in various ways to choose among alternative courses of

action or to select optimal strategies. Similarly, in designing computational systems that can reason, solve problems, and make decisions, we often try to design them to choose a course of action on the basis of some scale of measurement. Many times, these scales of measurement, whether used by human or artificial problem solvers, are based on subjective judgements. It is one of the goals of measurement theory to understand how humans can and should use subjective judgements to make better decisions. It is one of the goals of artificial intelligence to be able to handle highly complex problems by using subjective methods similar to those used by humans. In this project, we have studied properties of scales of measurement as they relate to their use by human or artificial decisionmakers in solving complex decisionmaking problems.

In many cases, the goal of a schedule is for items to be completed or to arrive at a given location by a certain time. For instance, AMC has developed a series of priorities or performance standards called UMMIPS for its schedules. Under UMMIPS, some highest priority items must reach their desired location within a short period of time, independent of cost, and there is a high "penalty" for not making the delivery on time. Lower priority items can arrive within a longer period of time and the penalty for missing the arrival time is lower. Mahadev, Pekec, and Roberts [1994a,b] introduced notions of desired arrival times, diverse performance standards, and varying penalties for missing desired arrival times, into the theory of scheduling. Building on these two papers, and motivated by these problems at AMC, paper [19] combines the results of these two earlier papers. It formulates precisely a variety of scheduling problems under performance constraints. In the problems analyzed, a number of items (equipment, people) have to be moved from an origin to a destination. It is assumed that each item has a desired arrival time at the destination and that we are penalized in some way for missing that time. The penalty can be applied only for a late arrival or, more generally, for both late and early arrivals, perhaps in a different way. It is assumed that we can only take a certain number of items from origin to destination each time that we schedule a trip (say because we have only a limited number of seats on each plane and only a limited number of planes). The goal is to minimize the total penalty. Paper [19] also considers the complication that the items have different priorities or status or importance. This complication is specifically motivated by the UMMIPS priorities. If there are different priorities, the penalty for early or late arrival can depend upon the priority. We make this problem precise, formulate a variety of specific penalty functions, and summarize a variety of relevant papers in the literature. We note that the introduction of priorities adds a complication if we take into account the way we measure them. Namely, scales of measurement often have certain arbitrary choices (such as of unit or zero point). If we allow admissible transformations of scale, we should ask if the optimal solution to the scheduling problem, the solution that minimizes the penalty, remains unchanged. Paper [19] observes that under some reasonable assumptions, it does not, and give conditions under which it does. At the beginning, the paper emphasizes analysis of the situation where the desired arrival time is the same for all items, and points out that, even here, we can be in the anomalous situation where an allowable change in the way we measure priorities changes the optimal schedule. These results have implications for how scheduling with priorities should be carried out. We then take the results one step further, analyzing scheduling problems in which not all items have the same desired arrival time. We give some general conditions under which a conclusion of optimality for a schedule is invariant under change of scale of the scale measuring priorities if the scale is an ordinal scale, one where all monotone increasing transformations of scale are admissible. In brief, these conditions require that the penalty increase

with increasing priority and with increasing distance from the desired arrival time, but that early and late arrivals be treated equally and that specifically the penalty involves a linear function of distance from desired arrival time. We note that the conclusion is false under certain relaxations of these conditions, such as asymmetric penalties for early and late arrival and quadratic functions of distance from desired arrival time. We show that the optimal solution to a variety of scheduling problems under performance constraints can be obtained by a simple greedy algorithm. We also present surprising examples to show that this greedy algorithm does not attain optimality in all situations.

Paper [20] studies a scheduling problem that is fundamental in scheduling theory, the problem of scheduling jobs on a single machine in which there are penalties for both late and early completions. Analogous to the results in paper [19], we point out here that if attention is paid to how certain parameters are measured, then a change of scale of measurement might lead to the anomalous situation where a schedule is optimal if these parameters are measured in one way, but not if they are measured in a different way that seems equally acceptable. We discuss conditions under which this anomaly is avoided. This paper was originally prepared under an earlier AFOSR project, but under the present grant we have substantially improved the results, extending them from the case of so-called interval scales to the case of so-called ordinal scales that are considerably more relevant to practical scheduling problems.

4.2. Taking Account of User Preferences for Schedules

As Keeney, et al. [1988] say in discussing the AI approach to advanced decision support and expert systems, "in today's approaches to user modelling the explicit representation of choices and, even more important, of preferences is usually missing. ... Consequently, the editors see the urgent need to integrate the deep knowledge of decision analysts into future systems of the type discussed here." One of the goals of this project has been to analyze the use of decision-analytic methods involving user choices and preferences in the construction of expert systems. We have brought preferences into analysis of scheduling problems through the use of graph coloring methods.

In studying the scheduling problem involving resource constraints (such as the constraint that two users cannot overlap in their scheduled times because they use the same resources or together would use more than the available resources), one notes that in its simplest form it is just the ordinary graph coloring problem: Assign colors (times or locations) to users so that if two users are related in a resource constraint, their corresponding colors (assigned times or locations) are different. However, in practical scheduling problems, users often specify a preferred time or location or sets of times or locations. For instance, a given military unit might specify one of a number of acceptable points of embarkation or one of a number of acceptable departure times. An acceptable schedule should be not only a graph coloring, but a graph coloring where the color assigned to x is in the set or list $R(x)$ of colors acceptable to x . Such a graph coloring is called a **list coloring** corresponding to $R(x)$. List colorings are difficult precisely because they involve multiple objectives and to find a list coloring, we must make trade-offs among these objectives. Thus, they model the kinds of issues raised by the Command Analysis Group at Air Mobility Command Headquarters [1994] in its emphasis on the need for methods to handle such trade-offs. In paper [22], we prove the well-known list coloring conjecture for line perfect graphs. We say that a graph is **k -choosable** if for every assignment of lists $R(x)$ of size k , there

is always a list coloring. The *choice number* of graph G , $ch(G)$, is the smallest k so that G is k -choosable. It is clear that the choice number is always at least as big as the chromatic number, and the difference between the two numbers can be arbitrarily large. However, when we define similar concepts for edge colorings, no one has discovered a graph for which the edge choice number is different from the edge chromatic number (chromatic index). The list coloring conjecture states that these last two numbers are equal. This conjecture can be restated as follows: The choice number of the line graph of H is always equal to the chromatic number of the line graph of H . Much of the work on choosability has been motivated by the Dinitz conjecture, which states that the list coloring conjecture holds for the complete bipartite graph $K(n,n)$. This conjecture was proved recently by Galvin. We have obtained a much more general result: The list coloring conjecture holds for all graphs whose line graph is perfect.

In practical scheduling problems in which individuals state their preferences through lists $R(x)$, it is usually impossible to satisfy everyone by assigning them each a color from their set $R(x)$. In paper [21] we have begun to develop a theory of list colorings where we accept a certain percentage of unsatisfied requests, i.e., where a certain percentage of the assignments give a color not in $R(x)$. We develop methods for finding list colorings that almost attain the desired conditions.

4.3. Conflicting Requests in Scheduling

We have already mentioned the importance to the Air Force of developing methods for dealing with trade-offs among multiple objectives. Such trade-offs also arise in scheduling in situations where we have conflicting requests. Suppose that user a and user b both wish an assignment at time or location x . Then there is a conflict (unless x has a large enough capacity to handle both, which is a special case which we have disregarded in this project). In general, one can study such conflicts by considering a bipartite digraph D whose vertices are elements of two sets, the set S of users and the set T of times or locations, and which has an arc from user a to time (location) x if user a is willing to depart at time (from location) x . Then there is a corresponding graph G whose vertices are the users and which has an edge between users a and b if and only if there are arcs from both a and b to the same x . There is a rather extensive literature devoted to the study of the graph G , which is called the *conflict graph* or *competition graph* corresponding to D .

This concept of conflict graph arises in a variety of applications. For instance, in communications, S is a set of transmitters and T a set of receivers and there is an arc in D from a in S to x in T if a message sent at a can be received at x ; the graph G represents conflict between transmitters. In coding, S is a set of codewords in a transmission alphabet, T a set of codewords in a receiving alphabet, and there is an arc in D from a in S to x in T if word a can be received as word x . The graph G represents confusability between codewords. Conflict graphs arise in ecology, where they are called competition graphs. Other applications arise in modelling of complex systems, and in particular in the analysis of models for such systems, in particular in the structural models, based on weighted digraphs or cross-impact matrices, that are used to study problems of energy, transportation, technology assessment, communications, and Naval manpower.

The *competition number* of a graph G is the smallest k so that G

together with k isolated vertices is a conflict/competition graph of an acyclic digraph. This number was introduced in 1978 by Roberts, who showed that the problem of its computation was equivalent to the problem of characterizing conflict/competition graphs of acyclic digraphs. Opsut [1982] showed that computation of this number was NP-complete. Based on an elimination algorithm developed by Parter and Rose for choosing the order of pivot points in Gaussian elimination, Roberts suggested in 1978 an elimination algorithm for computing the competition number. Opsut [1982] showed that this algorithm could overestimate the desired number. In paper [18], we have modified the elimination algorithm and showed that it correctly calculates the competition number for a large class of graphs. This paper, begun under an earlier AFOSR project, has been considerably revised, with significant improvement in the presentation of the algorithm and the justification that it works for a large class of graphs.

Using the original elimination algorithm, Roberts in 1978 found a formula for the competition number of a connected graph with no triangles, and this result has been widely used in the development of the theory of conflict/competition graphs. In paper [17], we have looked at the competition numbers of connected graphs with small numbers of triangles, and found exact solutions for the case where the graph has either one triangle or two triangles. This paper too was started under an earlier AFOSR project and has been considerably improved under the present project and the results strengthened.

The calculation of the competition number has led us to consider problems connected with cycle bases in graphs, as competition numbers sometimes can be calculated by finding cycle bases. In paper [23], we have obtained interesting relationships among different kinds of cycle bases, such as tree bases, face bases, triangle bases, induced bases, and ordering bases, and we present a number of useful results about counting bases.

Among other places, conflict graphs arise from command, control, and communications networks. One of the goals of our project has been to analyze conflict graphs of highly reliable routing networks analyzed in the literature. Among the candidates that have been widely studied as potentially highly efficient networks are those arising from circulant graphs and circulant matrices and their powers. In paper [2], we have studied powers of circulants. In particular, in bottleneck algebra (where addition and multiplication are replaced by the max and min operations), we consider the powers of a square matrix A . These powers are periodic, starting from a certain power A^k . The smallest such k is called the exponent of A and the length of the period is called the index of A . Cechlarova has characterized the matrices of index 1. We consider circulant matrices and determine when such matrices are idempotent (have exponent and period equal to 1). When the index is 1, we say that the circulant is strongly stable, and we show when this happens and observe that the result is equivalent to the result of Cechlarova for the case of circulant matrices. This paper too has been started under an earlier AFOSR project and significantly improved under this one.

4.4. Multiattribute Utility Theory

While the explicit representation of choices and of preferences is usually missing in decision support systems, a literature devoted to these topics in AI is beginning to be seen. In this project, we have briefly investigated the relevance of the theory of multiattribute utility functions to problems of scheduling and we

have investigated its uses in AI.

Our analysis has led us to consider some problems of multi-person game theory, where rewards depend upon complicated utility functions. The results are related to those already described in Section 3.3 in connection with knowledge discovery in databases. As we noted in Section 3.3, in paper [5], we note that some players may not like or know each other, so they cannot form a coalition. Let K be a fixed family of coalitions. The K -core is defined as the set of outcomes acceptable for all coalitions from K . The family K is called *stable* if the K -core is not empty for any normal form game. We prove that a family K of coalitions is stable if and only if K is a normal hypergraph.

An *effectivity function* is a Boolean function on the set IUA . An effectivity function is called *stable* if the core is not empty for any payoff function. The problem of characterizing stable effectivity functions seems, in general, very difficult. In paper [14], we apply a graph-theoretic approach to this problem. Using a graph based model, we obtain some necessary and sufficient conditions for stability in terms of perfect graphs, and we demonstrate that a conjecture by Berge and Duchet from 1983 is a special case of the considered problem of stability of effectivity functions.

References

Command Analysis Group at Air Mobility Command Headquarters, "XPY Research Directions," report, 1994.

De, S., "A Knowledge-Based Approach to Scheduling in an FMS," Annals of Operations Research, **12** (1988), 109-134.

Dechter, R., and Pearl, J., "Structure Identification in Relational Data," Artificial Intelligence, **58** (1992), 237-270.

Feldman, R., and Golumbic, M., "Interactive Scheduling as a Constraint Satisfiability Problem," Ann. Math. and Artif. Intell., **1** (1990), 49-73.

Hammer, P.L., and Kogan, A., "Optimal Compression of Propositional Horn Knowledge Bases: Complexity and Approximation," Artificial Intelligence, **64** (1993), 131-145.

Keeney, R.L., Möhring, R.H., Otway, H., Radermacher, F.J., and Richter, M.M. (eds.), Multi-attribute Decision Making Via O.R.-Based Expert Systems, Annals of Operations Research, Vol. 16, 1988.

Mahadev, N.V.R., Pekec, A., and Roberts, F.S., "On the Scheduling Problem with Priorities and Lateness/Earliness Penalties: Some Meaningful Conclusions with Nonconstant Arrival Times," Research Report RRR 20-94, RUTCOR, Rutgers University, June 1994. (a)

Mahadev, N.V.R., Pekec, A., and Roberts, F.S., "On the Scheduling Problem with Priorities and Lateness/Earliness Penalties: When is an Optimal Solution Not Optimal," Research Report RRR 19-94, RUTCOR, Rutgers University, June 1994. (b)

Maier, D., The Theory of Relational Databases, Computer Science Press, Rockville, MD, 1983.

Opsut, R.J., "On the Computation of the Competition Number of a Graph," SIAM J. Alg. & Discr. Meth., **3** (1982), 420-428.

Piatetsky-Shapiro, G., and Frawley, W.J., Knowledge Discovery in Databases, AAAI Press, 1991.

Ross, T.D., Axtell, M.L., Noviskey, M.J., and Gadd, D.A., "Pattern Theory for Systems Design," Midwest Symposium on Circuits and Systems, 1993.

Ross, T.D., Noviskey, M.J., Axtell, M.L., and Breen, M.A., "A Demonstration of a Robust Occam Based Learner," IEEE International Symposium on Information Theory, 1993.

RUTCOR

Grant Number F49620-95-1-0233

List of Publications Prepared under the Grant

Note: RRR means RUTCOR Research Report.

- [1] Anthony, Martin, "Accuracy of Techniques for the Logical Analysis of Data," RRR 23-96.
- [2] Atkin, A.O.L., Boros, E., Cechlarova, K., and Peled, U.N., "Powers of Circulants in Bottleneck Algebra," Linear Algebra and its Applications, in press.
- [3] Boros, E., and Cepek, O., "Perfect $0, +1$ Matrices," Annals of Discrete Math., to appear.
- [4] Boros, E., and Gurvich, V., "Stable Effectivity Functions and Perfect Graphs," preprint. (Revision of RRR 23-95)
- [5] Boros, E., Gurvich, V., and Vasin, A., "Stable Families of Coalitions and Normal Hypergraphs," Math. Soc. Sci, to appear.
- [6] Boros, E., Hammer, P.L., Ibaraki, T., and Kawakami, K., "Polynomial time Recognition of 2-Monotonic Positive Boolean Functions Given by an Oracle," SIAM J. on Computing, in press.
- [7] Boros, E., Hammer, P.L., Ibaraki, T., and Kogan, A., "Logical Analysis of Numerical Data," RUTCOR Technical Report RTR 3-95.
- [8] Boros, E., Hammer, P.L., Ibaraki, T., Kogan, A., Mayoraz, E., and Muchnik, I., "An Implementation of Logical Analysis of Data," RRR 22-96.
- [9] Boros, E., Ibaraki, T., and Makino, K., "Boolean Analysis of Incomplete Examples," SWAT '96, 1996, to appear. (See RRR 7-96.)
- [10] Boros, E., Ibaraki, T., and Makino, K., "Error-Free Fit and Best-Fit Extensions of Partially Defined Boolean Functions," preprint. (Revision of RRR 14-95)]
- [11] Boros, E., Ibaraki, T., and Makino, K., "Extensions of Partially Defined Boolean Functions with Missing Data," RRR 6-96.
- [12] Crama, Y., Ekin, O., and Hammer, P.L., "Variable Term Removal from Boolean Formulae," Discrete Applied Mathematics, to appear.
- [13] Ekin, O., Hammer, P.L., and Peled, U.N., "Horn Functions and Submodular Boolean Functions," Theoretical Computer Science, to appear.
- [14] Ekin, O., Hammer, P.L., Kogan, A., and Winter, P., "Distance-based Classification Methods," RRR 3-96.
- [15] Hammer, A., Hammer, P.L., and Muchnik, I., "Logical Analysis of Chinese

Productivity Patterns, RUTCOR Technical Report RTR 1-96.

[16] Hammer, P.L., Kogan, A., and Rothblum, U., "Relevance Functions and Evaluations of Sets of Variables of a Boolean Function," RRR 4-96.

[17] Kim, S., and Roberts, F.S., "Competition Number of Graphs with a Small Number of Triangles," preprint, December 1995. (Revision of RRR 9-95) (submitted for publication)

[18] Kim, S., and Roberts, F.S., "The Elimination Algorithm for the Competition Number," preprint, December 1995. (Revision of RRR 8-95) (submitted for publication)

[19] Mahadev, N.V.R., Pekec, A., and Roberts, F.S., "On an Aircraft Scheduling Problem with Priorities and Earliness/Tardiness Penalties," RRR 42-95.

[20] Mahadev, N.V.R., Pekec, A., and Roberts, F.S., "Single Machine Scheduling with Earliness and Tardiness Penalties: When is an Optimal Solution Not Optimal?" preprint. (Revision of RRR 28-95) (submitted for publication).

[21] Mahadev, N.V.R., and Roberts, F.S., "Almost Attainable List Colorings," in preparation.

[22] Peterson, D., "The List Coloring Conjecture for Perfect Line Graphs," in preparation.

[23] Xu, S., "Cycle Bases and their Applications," in preparation.

RUTCOR

Grant Number F49620-95-1-0233

Lectures Delivered

The following lectures on the topics of and related to the work of this project were delivered by the main participants:

Endre Boros

Invited seminar talk: "Structure of Horn functions,"
University of Tokyo, Japan, March 1995.

Invited seminar talk: "Perfect matrices,"
Tsukuba University, Japan, March 1995.

Invited seminar talk: "On the structure of Boolean functions,"
Kyoto University, Japan, March 1995.

Invited seminar talk: "Perfect graphs are kernel solvable,"
DIMACS, Rutgers University, May 1995.

Invited Plenary Lecture: "Graphs with no two cycles of equal length,"
International Conference on Hypergraphs and Symmetric Structures,
Balatonlelle, Hungary, June 1995.

Chairman of a session; invited talk: "Minimization of Horn functions,"
EURO XIV, Jerusalem, Israel, July 1995.

Organizer of a one day workshop on Boolean functions,
presentation: "Structure of Horn functions,"
EURO XIV, Jerusalem, Israel, July 1995.

Invited lecture: "Perfect 0,+1,-1 matrices,"
International Colloquium on Graph Theory,
Marseilles, France, September 1995.

Organizer and chairman of sessions, invited lecture:
"Structure of Horn Boolean functions,"
INFORMS, New Orleans, October 1995.

Series of invited (SIROCCO) lectures:
"Perfect graphs and kernel solvability,"
ARTEMIS, CNRS, Grenoble, France, November and December 1995.

Invited Seminar lecture: "Structure of Horn functions,"
IMAG, University of Grenoble, France, November 1995.

Invited Seminar lecture: "Minimization of Horn functions,"
Laboratory Leibniz, Institute National Polytechnique, Grenoble, France,
December 1995.

Plenary Lecture: "Binary Optimization,"
Optimization and its Applications, 2nd post-graduate workshop in
Operations Research, Ecole Polytechnique de Lausanne, Institute
National Polytechnique and University Joseph Fourier, Grenoble.
Seyssins, France, December 1995.

Invited Seminar talk: "Graphs and Games,"
Mathematical Institute, Hungarian Academy of Sciences,
Budapest, Hungary, December 1995.

Member of Program Committee; Invited lecture: "Structure of Horn rule
bases,"
4th International Conference on Mathematics and Artificial Intelligence,
Fort Lauderdale, Florida, January 1996.

Peter L. Hammer:

Invited talk.
International Meeting of INFORMS, Los Angeles, CA, April, 1995.

Minisymposium on Boolean Functions, Jerusalem, Israel, July, 1995.

14th European Conference of Operations Research, Jerusalem, Israel,
July, 1995, Organized three sessions on "Boolean Combinatorics and
Optimization".

Invited talk: "Horn and submodular Boolean functions".

Invited Talk.
5th International Symposium on Graph Theory and Combinatorics,
Marseille-Luminy, September, 1995.

International Meeting of INFORMS, New Orleans, LA, October 1995.
Presented tutorial on "Logical Analysis of Data".
Organized five sessions on "Boolean Functions".
Invited talks: "Horn functions and submodular Boolean functions"
and "Essential and redundant rules in Horn knowledge bases."

Alex Kogan:

Invited Talk: "Structure and Minimization of Horn Rule Bases"
INFORMS National Meeting, New Orleans, LA, October 1995.

Invited Talk: "Essential and Redundant Rules in Horn Knowledge Bases"
INFORMS National Meeting, New Orleans, LA, October 1995.

Fred S. Roberts:

Invited Talk at Session on Graph Theory: "Recent results about competition
graphs."
American Math Society meeting, Orlando, Florida, March 1995.

Departmental Colloquium, "Competition numbers and their applications."
Department of Mathematics, University of Louisville, March 1995.

Series of Plenary Talks, "Competition and conflict graphs."
First International Symposium on Combinatorics, Seoul, South Korea, August 1995.

Departmental Colloquium, "Mathematical modelling using graph theory."
Kyung-Hee University, Seoul, South Korea, August 1995.

Departmental Colloquium, "Applications of graph coloring."
Pohang Institute of Science and Technology, Pohang, South Korea, August 1995.

Plenary talk, "Applications of graph coloring."
Mathematical Association of America, NJ regional meeting, Cranford, NJ,
March 1996.

Departmental Colloquium, "Competition numbers and their applications"
Department of Mathematics, Dartmouth College, Hanover, NH, May 1996.

Plenary talk, "Role colorings and their applications."
International Conference on Graph Theory and Combinatorics, Kalamazoo,
MI, June 1996.

Invited talk, "Competition graphs"
International Colloquium on Combinatorics and Graph Theory, Balatonlelle,
Hungary, July 1996.

Invited talk at special session on mathematics and the social sciences:
"On the median procedure"
American Mathematical Society, summer national meeting, Seattle, WA,
August 1996.